



10/509527
Rec'd PCT/PTO 22 SEP 2004
PCT 2003 / 001244



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

REC'D 16 MAY 2003
WIPO PCT

#2

**PRIORITY
DOCUMENT**
SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 11 April 2003

An Executive Agency of the Department of Trade and Industry

BEST AVAILABLE COPY

The
Patent
Office

28MAR02 E707343-1 002776
P3/7700 1.09-02073435

THE PATENT OFFICE

A

28 MAR 2002

NEWPORT

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

1. Your reference

IP/P3304

2. Patent application number

(The Patent Office will fill in this part)

0207343.5

12 MAR 2002

3. Full name, address and postcode of the or of each applicant (underline all surnames)

QINETIQ LIMITED

Registered Office 85 Buckingham Gate
London SW1E 6PD
United Kingdom

Patents ADP number (if you know it)

8183857001

If the applicant is a corporate body, give the country/state of its incorporation

GB

4. Title of the invention

SIGNAL PROCESSING SYSTEM

5. Name of your agent (if you have one)

Bowdery Anthony Oliver

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

QINETIQ LIMITED
IP Formalities
A4 Bldg
Cody Technology Park
Ively Road
Farnborough
Hants GU14 0LX United Kingdom

Patents ADP number (if you know it)

8262792002

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number or earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an applicant, or
- c) any named applicant is a corporate body.

See note (d))

YES (b)

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form 0

Description 24

Claim(s) 6

Abstract 1

Drawing(s) 8 + 2

CF

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77) 1

Request for preliminary examination and search (Patents Form 9/77) 1

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11. I / We request the grant of a patent on the basis of this application.

Signature

D.J. RUSHTON

Date

27 March 2002

12. Name and daytime telephone number of person to contact in the United Kingdom

L BRUCKSHAW 01252 392722

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent of the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have attached 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

Signal Processing System

This invention relates to a system and method for processing signals to aid their classification and recognition. More specifically, the invention relates to
5 a modified process for training and using both Gaussian Mixture Models and Hidden Markov Models to improve classification performance, particularly but not exclusively with regard to speech.

Gaussian Mixture Models (GMMs) and Hidden Markov Models (HMMs) are
10 often used in signal classifiers to help identify an input signal when given a set of example inputs, known as training data. Uses of the technique include speech recognition, where the audio speech signal is digitised and input to the classifier, and the classifier attempts to generate from its vocabulary of words the set of words most likely to correspond to the input audio signal. Another
15 application is in radar, where radar signal returns from a scene are processed to provide an estimate of the contents of the scene. Published International specification WO02/08783 demonstrates the use of Hidden Markov Model processing of radar signals.

20 Before a GMM or HMM can be used to classify a signal, it must be trained with an appropriate set of training data to initialise parameters within the model to provide most efficient performance. There are thus two distinct stages associated with practical use of these models, the training stage and the classification stage. With both of these stages, data is presented to the
25 classifier in a similar manner. When applied to speech recognition, a set of vectors representing the speech signal is typically generated in the following manner. The incoming audio signal is digitised and divided into 10ms segments. The frequency spectrum of each segment is then taken, with windowing functions being employed if necessary to compensate for
30 truncation effects, to produce a spectral vector. Each element of the spectral vector typically measures the logarithm of the integrated power within each different frequency band. The audible frequency range is typically spanned by around 25 such contiguous bands, but one element of the spectral vector is conventionally reserved to measure the logarithm of the integrated power

across all frequency bands, i.e. the logarithm of the overall loudness of the sound. Thus, each spectral vector conventionally has around $25+1=26$ elements; in other words, the vector space is conventionally 26-dimensional. These spectral vectors are time-ordered and constitute the input to the HMM or GMM, as a spectrogram representation of the audio signal.

Training both the GMM and HMM involve establishing an optimised set of parameters associated with the processes using training data, such that optimal classification occurs when the model is subjected to unseen data.

- 10 A GMM is a model of the probability density function (PDF) of its input vectors (e.g. spectral vectors) in their vector space, parameterised as a weighted sum of Gaussian components, or classes. Available parameters for optimisation are the means and covariance matrices for each class, and prior class probabilities. The prior class probabilities are the weights of the weighted sum
- 15 of the classes. These adaptive parameters are typically optimised for a set of training data by an adaptive, iterative, re-estimation procedure such as the Expectation Maximisation (EM), and log-likelihood gradient ascent algorithms, which are well known procedures for finding a set of values for all the adaptive parameters that maximises the training-set average of the logarithm of the
- 20 model's likelihood function (log-likelihood). These iterative procedures refine the values of the adaptive parameters from one iteration to the next, starting from initial estimates, which may just be random numbers lying in sensible ranges.
- 25 Once the adaptive parameters of a GMM have been optimised, those trained parameters may subsequently be used for identifying the most likely of the set of alternative models for any observed spectral vector, i.e. for classification of the spectral vector. The classification step involves the conventional procedure for computing the likelihood that each component of the GMM
- 30 could have given rise to the observed spectral vector.

Whereas a GMM is a model of the PDF of individual input vectors irrespective of their mutual temporal correlations, a HMM is a model of the PDF of time-ordered *sequences* of input vectors. The adaptive parameters of an ordinary

HMM are the observation probabilities (the PDF of input vectors given each possible hidden state of the Markov chain) and the transition probabilities (the set of probabilities that the Markov chain will make a transition between each pair-wise combination of possible hidden states).

5

For the case of an ordinary GMM-based HMM, the observation probabilities are parameterised as a weighted sum of Gaussian components ('classes'), i.e. the observation probabilities are parameterised as GMMs. Thus, a prescription for optimising the HMM's observation probabilities can be re-cast as a prescription for optimising the associated GMM's class means, covariance matrices and prior class probabilities.

10

Training, or optimisation, of the adaptive parameters of a HMM is done so as to maximise the overall likelihood function of the model of the input signal, such as a speech sequence. One common way of doing this is to use the Baum-Welch re-estimation algorithm, which is a development of the technique of expectation maximisation of the model's log-likelihood function, extended to allow for the probabilistic dependence of the hidden states on their earlier values in the speech sequence. A HMM is first initialised with initial, possibly random, assumptions for the values of the transition and observation probabilities.

15

20

For each one of a set of sequences of input training vectors, such as speech-sequences, the Baum-Welch forward-backward algorithm is applied, to deduce the probability that the HMM could have given rise to the observed sequence. On the basis of all these per-sequence model likelihoods, the Baum-Welch re-estimation formula updates the model's assumed values for the transition probabilities and the observation probabilities (i.e. the GMM class means, covariance matrices and prior class probabilities), so as to maximise the increase in the model's average log-likelihood. This process is iterated, using the Baum-Welch forward-backward algorithm to deduce revised model likelihoods for each training speech-sequence and, on the basis of these, using the Baum-Welch re-estimation formula to provide further updates to the adaptive parameters.

25

30

Each iteration of the conventional Baum-Welch re-estimation procedure can be broken down into five steps for every GMM-based HMM: (a) applying the Baum-Welch forward-backward algorithm on every training speech-sequence,
5 (b) the determination of what the updated values of the GMM class means should be for the next iteration, (c) the determination of what the updated values of the GMM class covariance matrices should be for the next iteration, (d) the determination of what the updated values of the GMM prior class probabilities should be for the next iteration, and (e) the determination of what
10 the updated values of the HMM transition probabilities should be for the next iteration. Thus, the Baum-Welch re-estimation procedure for optimising a GMM-based HMM can be thought of as a generalisation of the EM algorithm for optimising a GMM, but with the updated transition probabilities as an extra, fourth output.

15 For certain applications, HMMs are employed that do not have their observation probabilities parameterised as GMMs, but instead use lower level HMMs. Thus, a hierarchy is formed that comprises at the top a "high level" HMM, and at the bottom a GMM, with each layer having its observation
20 probabilities defined by the next stage down. This technique is common in subword-unit based speech recognition systems, where the structure comprises two nested levels of HMM, with the lowest one having GMM based observation probabilities.

25 The procedure for optimising the observation probabilities of a high-level HMM reduces to the conventional procedure for optimising both the transition probabilities and the observation probabilities (i.e. the GMM parameters) of the ordinary HMMs at the lower level, which is as described above. The procedure for optimising the high-level HMM's transition probabilities is the
30 same as the conventional procedure for optimising ordinary HMMs' transition probabilities, which is as described above.

HMMs can be stacked into multiple-level hierarchies in this way. The procedure for optimising the observation probabilities at any level reduces to

the conventional procedure for optimising the transition probabilities at all lower levels combined with the conventional procedure for optimising the GMM parameters at the lowest level. The procedure for optimising the transition probabilities at any level is the same as the conventional procedure for optimising ordinary HMMs' transition probabilities. Thus, the procedure for optimising hierarchical HMMs can be described in terms of recursive application of the conventional procedures for optimising the transition and observation probabilities of ordinary HMMs.

- 10 Once the HMM's adaptive parameters have been optimised, the trained HMM may subsequently be used for identifying the most likely of a set of alternative models of an observed sequence of input vectors - spectral vectors in the case of speech classification. This process conventionally is achieved using the Baum-Welch forward-backward algorithm, which computes the likelihood of generating the observed sequence of input vectors from each of a set of alternative HMMs with different optimised transition and observation probabilities.

- The classification methods described above have certain disadvantages.
- 20 When optimising the observation probabilities of the GMMs, and hence of the HMMs that may be hierarchically above them, as well as the transition probabilities of the HMM, there is a tendency for the optimisation to get caught in local minima, which prevents the system from achieving optimal classification. This can often be attributed to a tendency for class likelihood-PDFs to become "tangled up" with one another if they are free to become too highly anisotropic. Also, regarding speech recogniser technology, current recognisers are poor at capturing subtle variations and intrinsic characteristics of real speech, such as the full, specific variability of speakers' vowels under very different speaking conditions. In particular, individual vowels occupy complex shapes in spectral vector space, and attempting to represent these shapes as Gaussian distributions, as is conventionally done, can lead to unfaithful representation of the speech sounds.

According to the present invention there is provided a signal processing system for processing a plurality of multi-element data encoding vectors, the system:

- 5 - having means for deriving the data encoding vectors from input signals;
- being arranged to process the data encoding vectors using a Gaussian Mixture Model (GMM) based Hidden Markov Model (HMM), the GMM based HMM having at least one class mean vector having multiple elements;
- 10 - being arranged to process the elements of the class mean vector(s) by an iterative optimisation procedure;
- characterised in that the system is also arranged to scale the elements of the class mean vector(s) during the optimisation procedure to provide for the class mean vector(s) to have constant modulus at each iteration, and to
- 15 normalise the data encoding vectors input to the GMM based HMM.

A GMM-based HMM is a generalisation of a GMM such that the HMM has observation probabilities parameterised as Gaussian PDFs or weighted sums of Gaussian PDFs, i.e. as a GMM. The observation probabilities of a GMM-based HMM are parameterised as a GMM, but the GMM-based HMM is not

20 itself a GMM. An input stage can be added to a GMM based HMM however, where this input stage comprises a simple GMM. The log-likelihood of a GMM-based HMM is the log-likelihood of an HMM whose observation probabilities are constrained to be parameterised as GMMs; it is not the log-likelihood of a GMM. Consequently, the optimisation procedure of a GMM-based HMM is not the same as that of a GMM.

Preferably the moduli of the mean vectors of each of the GMMs are rescaled after each iteration of the optimisation procedure so that they are all of equal value.

30 Most signal processing systems of the type discussed in this specification incorporate a GMM that represents the probability density function of all data encoding vectors in the training sequence. The constraint of limiting the elements of the class mean vector to have constant modulus leads to

simplified processing of the GMMs making up the signal processing system, as the class means of each GMM will lie on the surface of a hypersphere having dimensionality $(n - 1)$, where n is the dimension of an individual vector.

- 5 Preferably each covariance matrix is constrained so as to be isotropic and diagonal, and to have a variance constrained to be a constant value. This eliminates the possibility of certain classes of severe local minima associated with highly anisotropic Gaussian components, and so prevents such sub-optimal configurations from forming during the training process. Note that a
- 10 covariance matrix that is so constrained may be regarded mathematically as a scalar value, and hence a scalar value may be used to represent such a covariance matrix.

- Each GMM, and therefore GMM based HMM, has a set of prior class
- 15 probabilities. Preferably the prior class probabilities associated with the GMM are constrained to be equal, and to remain constant throughout the optimisation procedure.

- Prior art signal processing systems incorporating GMMs generally avoid
- 20 putting constraints on the model parameters; other than that covariance matrices are on occasion constrained to be equal across classes, requirements are rarely imposed on the class means, covariance matrices, prior class probabilities and hidden-state transition probabilities other than that their values are chosen to make the average log-likelihood as large as
- 25 possible.

- Preferably, each data encoding vector that is also an input vector, derived from the input signal during both training and classifying stages of using the GMM is constrained such that its elements x_i are proportional to the square
- 30 roots of the integrated power within different frequency bands. Advantageously, the elements of each such data encoding vector are scaled such that the squares of the elements of the vector sum to a constant value that is independent of the total power of the original input.

Preferably each such data encoding vector is augmented with the addition of one or more elements representing the overall power in the vector. The scaling of the vector elements described above removes any indication of the power, so the additional element(s) provide the only indication of the power, or
5 loudness, within the vector. Clearly, the computation of the value of the elements representing power would need to be based on pre-scaled elements of the vector.

Certain applications, notably subword-unit based models advantageously
10 employ a HMM that uses as its observation probability a GMM constrained according to the current invention, that likewise acts as the observation probability for a further HMM. In this way, a hierarchy of HMMs can be built up, in the manner of the prior art, but with the difference that the constraints on the model parameters according to the current invention are applied at
15 each level of the hierarchy.

Advantageously, the hierarchy may incorporate two GMMs as two lower levels, with a HMM at the highest level. The lowest level GMM provides posterior probabilities as a data encoding vector to a second, higher level
20 GMM. This second GMM provides observation probabilities to a HMM at the third level. This arrangement allows individual speech-sounds to be represented in the spectral-vector space not as individual Gaussian ellipsoids, as is conventional, but as assemblies of many smaller Gaussian hypercircles tiling the unit hypersphere, offering in the potential for more faithful
25 representation of highly complex-shaped speech-sounds, and thus improved classification performance.

Note that in this specification the terms "input vector" and "spectral vector" are used interchangeably in the context of providing an input to the lowest level of
30 the system hierarchy. The vector at this level may represent the actual power spectrum of the input signal, and hence be spectral coefficients, or may represent some modified form of the power spectrum. In practice, the input vector will generally represent a power spectrum of a segment of a temporal input signal, but this will not be the case for all applications. Further

processing of the temporal input signal is used in some applications, e.g. cosine transform. A "data encoding vector" is, within this specification, any vector that is used as an input to any level of the hierarchy, depending on the context, i.e. any vector that is used as the direct input to the particular level of the hierarchy being discussed in that context. A data encoding vector is thus
5 an input vector only when it represents the information entering the system at the lowest level of the hierarchy.

Note also that normalising a vector is the process of rescaling all its elements
10 by the same factor, in order to achieve some criterion defined on the whole vector of elements. What that factor is depends on the criterion chosen for normalisation. A vector can generally be normalised by one of two useful criteria; one is to normalise such that the elements sum to a constant after normalisation, the other is to normalise such that the squares of the elements
15 sum to a constant after normalisation. By the first criterion, the rescaling factor should be proportional to the reciprocal of the sum of the values of the elements before normalisation. By the second criterion, the rescaling factor should be proportional to the reciprocal of the square root of the sum of the squares of the values of the elements before normalisation. A vector of
20 exclusive probabilities is an example of a vector normalised by the first criterion, such that the sum of those probabilities is 1. A (real-valued) unit vector is an example of a vector normalised according to the second criterion; the sum of the squares of the elements of a (real-valued) unit vector is 1. A vector whose elements comprise the square roots of a set of exclusive
25 probabilities is also an example of a vector normalised by the second criterion.

According to another aspect of the current invention there is provided a method of processing a signal, the signal comprising a plurality of multi-element data encoding vectors, wherein the data encoding vectors are
30 derived from an analogue or digital input, and where the method employs at least one Gaussian Mixture Model (GMM) based Hidden Markov Model (HMM), the GMM based HMM having at least one class mean vector having multiple elements, and the elements of the class mean vector(s) are optimised in an iterative procedure, characterised in that the elements of the class mean

vectors are scaled during the optimisation procedure such that the class mean vectors have a constant modulus at each iteration, and the data encoding vectors input to the GMM based HMM are processed such that they are normalised.

5

Note that the user(s) of a system trained according to the method of the current invention may be different to the user(s) who performed the training. This is due to the distinction between the training and the classification modes of the invention

10

According to another aspect of the current invention there is provided a computer program designed to run on a computer and arranged to implement a signal processing system for processing one or more multi-element input vectors, the system:

15

- having means for deriving the data encoding vectors from input signals;

- being arranged to process the data encoding vectors using a Gaussian Mixture Model (GMM) based Hidden Markov Model (HMM), the GMM based HMM having at least one class mean vector having multiple

20

- elements;
- being arranged to process the elements of the class mean vector(s) by an iterative optimisation procedure;

characterised in that the system is also arranged to scale the elements of the class mean vector(s) during the optimisation procedure to provide for the

25

class mean vector(s) to have constant modulus at each iteration, and to normalise the data encoding vectors input to the GMM based HMM.

30

The present invention can be implemented on a conventional computer system. A computer can be programmed to so as to implement a signal processing system according to the current invention to run on the computer hardware.

According to another aspect of the current invention there is provided a speech recogniser incorporating a signal processing system for processing one or more multi-element input vectors, the recogniser:

- having means for deriving the data encoding vectors from input signals;
- being arranged to process the data encoding vectors using a Gaussian Mixture Model (GMM) based Hidden Markov Model (HMM), the GMM based HMM having at least one class mean vector having multiple elements;
- being arranged to process the elements of the class mean vector(s) by an iterative optimisation procedure; characterised in that the system is also arranged to scale the elements of the class mean vector(s) during the optimisation procedure to provide for the class mean vector(s) to have constant modulus at each iteration, and to normalise the data encoding vectors input to the GMM based HMM.

A speech recogniser may advantageously incorporate a signal processing system as described herein, and may incorporate a method of signal processing as described herein.

The current invention will now be described in more detail, by way of example only, with reference to the accompanying Figures, of which:

Figure 1 diagrammatically illustrates a typical hardware arrangement suitable for use with the current invention when implemented in a speech recogniser.

Figure 2 shows in block diagrammatic form the conventional re-estimation procedure adopted by the prior art systems employing GMM or HMM based classifiers;

Figure 3 shows in block diagrammatic form one of the pre-processing stages carried out on input vectors based on frames of speech, relating to the frame's spectral shape;

Figure 4 shows in block diagrammatic form a further pre-processing stage carried out on the input vectors relating to the overall loudness of a frame of speech;

5

Figure 5 shows in block diagrammatic form the modified re-estimation procedure of GMMs or ordinary, or hierarchical HMMs as per the current invention;

10 Figure 6 shows in more detail the class mean re-scaling constraint shown in Figure 5;

Figure 7 shows in block diagrammatic form the implementation of a complete system; and

15

Figure 8 shows graphically one advantage of the current invention using the example of a simplified three dimensional input vector space.

The current invention would typically be implemented on a computer system
20 having some sort of analogue input, an analogue to digital converter, and digital processing means. The digital processing means would comprise a digital store and a processor. As shown in Figure 1, a speech recogniser embodiment typically has a microphone 1 acting as a transducer from the speech itself, the electrical output of which is fed to an analogue to digital
25 converter (ADC) 2. There may also be some analogue processing before the ADC (not shown). The ADC feeds its output to a circuit 3 that divides the digital signal into 10ms slices, and carries out a spectral analysis on each slice, to produce a spectral vector. These spectral vectors are then fed into the signal processor 4, in which is implemented the current invention. The
30 signal processor 4 will have associated with it a digital storage 5. Some applications may have as an input a signal that has been digitised at some remote point, and so wouldn't have the ADC. Other hardware configurations are also possible within the scope of the current invention.

A typical signal processing system of the current invention will comprise a simple GMM and a GMM-based HMM, together used to classify an input signal. Before either of those models can be used for classification purposes, they must first be optimised, or trained, using a set of training data. There are
5 thus two distinct modes of operation of a classification model: the training phase, and the classification phase.

Figure 2 shows generically the steps used by prior art systems in training both a GMM and a HMM based classifier. Figure 2 depicts the optimisation of
10 hierarchical GMM-based HMMs as well as the optimisation of ordinary GMM-based HMMs and simple GMMs, because the steps relating to initialising and re-estimating HMM transition probabilities relate to the initialisation and re-estimation of HMM transition probabilities at all levels of the hierarchy. The flow chart is entered from the top when it is required to establish an
15 improved set of parameters in the model to improve the classification performance. First the various classes need to be initialised, these being the class means, class covariance matrices and prior class probabilities. HMMs have the additional step of initialising the transition probabilities. These initialisation values may be random, or they may be a "best guess" resulting
20 either from some previous estimation procedure or from some other method.

These initialisations form the adaptive parameters for the first iteration of the training procedure, which proceeds as follows. An data encoding vector or vector sequence (for the HMM case) from the training sequence is obtained,
25 and processed using a known re-estimation procedure. For GMMs the EM algorithm is often used, and for HMMs the Baum-Welch re-estimation procedure is commonplace. This is the inner loop of the re-estimation procedure, and is carried out for all data encoding vectors in the training sequence.

30 Following this, the information gained during the inner loop processing is used to compute the new classes and, for the HMM case, the new transition probabilities. Convergence of this new data is tested by comparing it with the previous set or by judging whether the likelihood function has achieved a

stable minimum, and the process re-iterated if necessary using the newly computed data as a starting point.

5 Moving to the current invention, one embodiment of the current invention applied to speech recognition employs a modified spectral vector that is pre-processed in a manner that is different from the conventional log-power representation of the prior art. The spectral vector itself comprises a spectral representation of a 10ms slice of speech, divided up into typically 25 frequency bins.

10

The objective of the first stage of the pre-processing is that elements x_i ($i=1, \dots, m$) of the n -dimensional ($m \leq n$) spectral vector \mathbf{x} should be proportional to the square roots $\sqrt{P_i}$ of integrated power P_i within different frequency bands, rather than the conventional logarithms of integrated power within different frequency bands. Further, elements x_i ($i=1, \dots, m$) should be scaled such that their squares should sum to a constant A that is independent of the total power integrated across all frequency bands within the frame corresponding to that spectral vector. Thus, if the frame is sampled into m frequency bands, m of the elements x_i of the n -dimensional ($m \leq n$) spectral vector \mathbf{x} should satisfy

20

$$x_i = A \sqrt{P_i} / \sum_{j=1}^m P_j \quad (i=1, \dots, m) \quad (\text{Equation 1})$$

which implies $\sum_{j=1}^m x_j^2 = A$.

The value of the constant A has no functional significance; all that matters is that it doesn't change from one spectral vector to the next.

25

The advantage of this normalised square root power representation for spectral vectors is that the degree of match of the shape of spectral vector x_i ($i=1, \dots, m$), compared with a class mean vector w_i ($i=1, \dots, n$), is then proportional to the scalar product $\sum_{i=1}^m x_i w_i$, irrespective of the modulus (vector length) of the template. This provides the freedom to constrain the modulus of the template without losing the functionality of being able to

30

determine the degree of match of the template by computing the scalar product.

The steps involved in the novel encoding of spectral vectors are represented in the flow diagram of Figure 3 and listed as follows (a-e). After (a) choosing a value for the constant A to be used for all frames of speech, (b) the first step to be applied for each individual frame of speech is the same as the conventional process for conducting a spectral analysis in order to obtain m values of the integrated power P_i ($i=1, \dots, m$) within m different frequency bands spanning the audible frequency range. Then, instead of taking the logarithms of these power-values as is conventional in the prior art, (c) their sum $\sum_{j=1}^m P_j$ and (d) their square roots $\sqrt{P_i}$ ($i=1, \dots, m$) are computed. (e) each square-root value $\sqrt{P_i}$ is then divided by the total power $\sum_{j=1}^m P_j$ (and multiplied by a constant scaling factor A as desired) to obtain elements x_i ($i=1, \dots, m$) of the novel encoding of the spectral vector defined by equation 1.

As a second part of the pre-processing of the spectral vectors, the vector is also augmented with the addition of extra elements that represent the overall loudness of the speech at that frame, i.e. the total power $\sum_{j=1}^m P_j$ integrated across all frequency bands.

This is particularly useful in conjunction with the novel way of encoding spectral shape defined by equation 1. This is because elements x_i ($i=1, \dots, m$) defined by equation 1 are clearly independent of the overall loudness $\sum_{j=1}^m P_j$ and therefore encode no information about it, so those m elements need to be augmented with additional information if the spectral vector is to convey loudness information.

In the current embodiment, two extra elements x_{m+1} and x_{m+2} are added to the spectral vector, beyond the m elements used to encode the spectral shape. Thus the spectral vector will have $n = m+2$ dimensions. These two elements depend on the overall loudness $L \equiv \sum_{j=1}^m P_j$ in the following way:

$$x_{m+1} = B \frac{f(L)}{\sqrt{[f(L)]^2 + [g(L)]^2}}, \quad x_{m+2} = B \frac{g(L)}{\sqrt{[f(L)]^2 + [g(L)]^2}} \quad (\text{Equation 2})$$

where $f()$ and $g()$ are two (different) functions of the overall loudness L , and B is a constant. The significance of B is that the ratio B/A determines the relative contributions to the squared modulus $|\mathbf{x}|^2 = \mathbf{x} \cdot \mathbf{x} = \sum_{j=1}^n x_j^2$ made by

5 the two subsets of elements ($i = m+1, m+2$) and ($i = 1, \dots, m$); the values of these contributions are clearly B^2 and A^2 respectively. The ratio B/A may therefore be used to control the relative importance assigned to overall loudness and spectral shape in the coding of spectral vectors; for example, choosing $B = 0$ assigns no importance to overall loudness, while choosing
10 similar values of A and B assigns similar importance to both aspects of the speech. The value of $A^2 + B^2$ can be chosen to be 1 for simplicity, which will make the squared modulus $|\mathbf{x}|^2 = \mathbf{x} \cdot \mathbf{x} = \sum_{j=1}^n x_j^2 = A^2 + B^2$ equal to 1 for all spectral vectors regardless of their speech content.

The advantages of this novel representation of loudness are (a) that the
15 moduli of all spectral vectors will have the same constant value regardless of overall loudness, which frees one to constrain the moduli of templates (class means) $\mathbf{w} = (w_1, \dots, w_n)$, as is proposed in the main claims, and (b) that the ratio B/A may be used to control the relative importance assigned to overall loudness and spectral shape in the coding of spectral vectors.

20 Possible choices for the functions $f()$ and $g()$ include

$$f(L) = \sin\left(\frac{\pi}{2} \frac{\log L - \log L^{\min}}{\log L^{\max} - \log L^{\min}}\right), \quad g(L) = \cos\left(\frac{\pi}{2} \frac{\log L - \log L^{\min}}{\log L^{\max} - \log L^{\min}}\right) \quad (\text{Equation 3})$$

where L^{\min} and L^{\max} are constants chosen to correspond to the quietest and loudest volumes (total integrated power) typically encountered in individual frames of speech.

25 Useful values for the pair of constants (A, B) are $(1, 0)$, $(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$ and $(\sqrt{\frac{2}{n}}, \sqrt{\frac{m}{n}})$, which all satisfy $A^2 + B^2 = 1$.

Once functions $f()$ and $g()$ and constants B , L^{\min} and L^{\max} , to be used for all frames of speech, have been chosen, the steps involved in the process
30 required to incorporate the loudness encoding as described above are shown

in Figure 4. The process involves (a) summing the integrated powers P_i within m frequency ranges $i=1,\dots,m$ for each frame of speech to obtain the overall loudness L for that frame of speech, (b) evaluating the two extra elements x_{m+1} and x_{m+2} for that frame of speech according to equation 2, and
5 (c) for that frame of speech appending the two extra elements to the m elements obtained from the process of figure 4 to obtain an $n=m+2$ dimensional spectral vector incorporating the novel encodings of spectral shape and loudness.

10 The steps as shown in Figures 3 and 4 comprise the pre-processing of the spectral vectors according to the embodiment of the current invention.

The input vectors pre-processed as described above are used when optimising the various parameters of the GMMs and GMM-based HMMs. The
15 inner loop of the optimisation procedure, as described in relation to Figure 1 above, is done using convention methods such as EM re-estimation and Baum-Welch re-estimation, respectively. Further novel stages are concerned with applying constraints to the parameters in between iterations of this inner loop.

20 Figure 5 shows the re-estimation procedure of the current invention, with additional processes present as compared to that shown in Figure 2. These additional processes relate to the initialisation of the classes before the iterative part of the procedure starts, and to the rescaling of the class means
25 following each iteration to take into account the constraints to be imposed. Note that for the HMM case the transition probability processing is unchanged from the prior art.

One of the constraints is concerned with the class mean vectors of the GMM
30 or HMM. The constraint takes the form of re-scaling the set of n -dimensional vectors $w_j = (w_{j1}, \dots, w_{jn})$ which represent the class means.

This constraint is applied to all the class means, as soon as they have been re-estimated, every time they are re-estimated (by the EM or Baum-Welch re-

estimation procedures for example), and also when they are first initialised (see Figure 5). These extra steps, illustrated in the flow diagram of Figure 5, are (a) by summing the squares of its elements and then taking the square root of the sum, the modulus $|w_j|$ of each of the N re-estimated class means w_j is first computed as

$$|w_j| = \sqrt{\sum_{i=1}^n w_{ji}^2} \quad (\text{Equation 4})$$

for all N classes $j = 1, \dots, N$; (b) after computing the modulus $|w_j|$ of each re-estimated class mean, all the elements of each class mean are divided by that corresponding modulus, i.e.

$$w_{ji} \rightarrow D \frac{w_{ji}}{|w_j|}, \text{ for all elements } i = 1, \dots, n \text{ of all GMM classes } j = 1, \dots, N$$

(Equation 5)

These steps have the effect of re-scaling all the class means w_j to constant modulus D until the next iteration of their re-estimation, after which they are re-scaled again to constant modulus D by applying these steps again, as depicted in Figure 5. The value of the constant D is preferably set equal to the modulus $|x|$ of the data vectors x . (For example, for a GMM receiving input data having moduli $|x| = \sqrt{A^2 + B^2}$, the value of D should be set equal to $\sqrt{A^2 + B^2}$.)

The advantages of re-scaling the class means to constant modulus are that this encourages speech recognition algorithms to adopt novel encodings of speech data that may improve speech classification performance (such as hierarchical sparse coding), and that it may reduce the vulnerability of speech recognition algorithms to becoming trapped in undesirable sub-optimal configurations ('local minima') during training. These advantages result from the fact that the dynamics of learning have simplified degrees of freedom because the class means are constrained to remain on a hypersphere (of radius D) as they adapt.

Re-scaling class means w_j to constant modulus is particularly appropriate in conjunction with scaling data vectors x to constant modulus. This is because

the degree of match between a data vector x and a class mean w_j can then be determined purely by computing the scalar product $w_j \cdot x$.

Further to the embodiment of the current invention, the covariance matrices C_j of the Gaussian distributions that constitute the GMMs are constrained to be isotropic and of constrained variance V , i.e. they are not optimised according to the conventional re-estimation procedures for covariance matrices (such as the EM algorithm for GMMs and the Baum-Welch procedure for GMM-based HMMs), but are defined once and for all in terms of the isotropic Identity Matrix I and the constrained variance V by

$$C_j \equiv VI \text{ for all classes } j = 1, \dots, N \quad (\text{Equation 6})$$

V is a free parameter chosen (for example by trial and error) to give the speech recognition system best classification performance; V must be greater than zero, as a covariance matrix has non-negative eigenvalues, and V is preferably significantly smaller than the value of D^2 . The benefit of setting V much smaller than D^2 is that it leads to a sparse distribution of the first level simple GMM's posterior probabilities, which in the main embodiment feed the data encoding vector space of the GMM-based HMM at the second level. This is because each Gaussian component of the first level simple GMM will individually only span a small area on the spectral vector hypersphere.

This process for choosing covariance matrices involves the following steps: (a) choosing a value for the constant of proportionality V so as to optimise the classification performance, for example by trial and error, (b) setting all the diagonal elements of the class covariance matrices equal to V , and (c) setting all the off-diagonal elements of the class covariance matrices equal to zero. Thus, the covariance matrix according to this embodiment of the present invention is both isotropic and diagonal.

Used in conjunction with the above techniques for constraining the moduli of data vectors x and class means w_j , constraining the class covariances in this way gives the advantage of encouraging speech recognition algorithms to adopt novel encodings of speech data that may improve speech recognition

performance (such as hierarchical sparse coding), and reducing the vulnerability of speech recognition algorithms to becoming trapped in undesirable sub-optimal configurations ('local minima') during training. Sparse coding results from representing individual speech-sounds as assemblies of many small isotropic Gaussian hypercircles tiling the unit hypersphere in the spectral-vector space, offering in the potential for more faithful representation of highly complex-shaped speech-sounds than is permitted by representation as a single anisotropic ellipsoid, and thus improved classification performance.

Because this constraint does away with the need for the conventional unconstrained re-estimation of the covariance matrices, Figure 5's modified procedure for optimising GMMs does not involve re-estimation of covariance matrices as does the conventional procedure of Figure 2.

A further constraint imposed on this embodiment of the current invention relates to the choice of prior class probabilities. The N prior probabilities $\text{Pr}(j)$ for the GMM classes $j = 1, \dots, N$ may be constrained to be constants, i.e. not optimised according to the conventional re-estimation procedures for prior class probabilities (such as the EM algorithm for GMMs and the Baum-Welch procedure for GMM-based HMMs), but are defined once and for all by the step of setting

$$\text{Pr}(j) = 1/N \text{ for all classes } j = 1, \dots, N \quad (\text{Equation 7})$$

Used in conjunction with the above innovations for constraining the moduli of data vectors \mathbf{x} , class means \mathbf{w}_j and the covariance matrices \mathbf{C}_j , constraining the prior class probabilities in this way gives the advantage of reducing the vulnerability of speech recognition algorithms to becoming trapped in undesirable sub-optimal configurations ('local minima') during training.

Because this innovation does away with the need for the conventional unconstrained re-estimation of the prior class probabilities, Figure 5's modified procedure for optimising GMMs does not involve re-estimation of prior class probabilities as does the conventional procedure of Figure 2.

It will be understood by people skilled in the relevant arts that the constraints applied to a GMM or HMM as described above in the training phase of the model will equally need to be applied during the classifying phase of use of the models. If they were employed during training, the steps for encoding spectral shape and overall loudness according to the present invention as described above will need to be applied to every spectral vector of any new speech to be classified.

An implementation of the invention, which combines all of the constraints detailed above, is illustrated in Figure 6. This implementation uses conventional spectral analysis of each frame of speech, followed by novel steps described above to encode both spectral shape and overall loudness into each spectral vector and to scale every spectral vector's modulus to the constant value of 1. The parameters A and B are both set to equal $1/\sqrt{2}$ and D is set equal to 1.

Such unit-modulus spectral vectors are input to a GMM having a hundred Gaussian classes ($N = 100$), with class means all constrained to have moduli equal to 1, with class prior probabilities all constrained to have constant and equal values of $1/100$, and covariance matrices constrained to be isotropic and to have constant variances (i.e. not re-estimated at each iteration according to a procedure such as the EM algorithm). A good choice for that constant variance V has been found to be 0.01, although other values could be chosen by trial and error so as to give best speech classification performance of the whole system; the right choice for V will lie between 0 and 1. For each spectral vector input to this GMM, posterior probabilities for the classes are computed in the conventional way.

Each set of GMM posterior probabilities computed above for each spectral vector are used to compute unit-modulus data-encoding vectors for input to an ordinary GMM-based HMM by taking the square roots of those posterior probabilities.

These unit-modulus data-encoding vectors are input to the HMM as observation vectors. The class means of the Gaussian mixture that constitutes

the parameterisation of the HMM's observation probabilities are, all constrained to have moduli equal to 1. The number N of Gaussian classes used to parameterise the HMM's observation probabilities is chosen by trial and error so as to give best speech classification performance of the whole system. The prior probabilities of those classes are then determined by that choice of N ; they are all constrained and set equal to $1/N$. The covariance matrices of those classes are all constrained to be isotropic and to have constant variances (i.e. not re-estimated unconstrained according to a procedure such as the EM algorithm). The choice of that constant variance V would be determined by trial and error so as to give best speech classification performance of the whole system; the right choice for V will lie between 0 and 1.

The preferred implementation of the invention can be operated in training mode and classification mode. In classification mode, the HMM is used to classify the input observation vectors according to a conventional HMM classification method (Baum-Welch forward-backward algorithm or Viterbi algorithm), subject to the modifications described above.

In training mode, (a) the GMM is optimised for the training of unit-modulus spectral vectors (encoded as described above) according to a conventional procedure for optimising GMM class means (e.g. the EM re-estimation algorithm), subject to the innovative modifications to re-scale the GMM class means to have constant moduli equal to 1, and to omit the conventional steps for re-estimating the GMM class covariance matrices and prior class probabilities. (b) Once the GMM has been optimised, it is used as described above to compute a set of data-encoding vectors from the training set of speech spectral vectors. (c) This set of data-encoding vectors is then used for training the HMM according to a conventional procedure for optimising HMM class means (e.g. the Baum-Welch re-estimation procedure), subject to the innovative modifications to re-scale the HMM class means to have constant moduli equal to 1, and to omit the conventional steps for re-estimating the HMM class covariance matrices and prior class probabilities. No modification is made to the conventional steps for re-estimating HMM transition

probabilities; the conventional Baum-Welch re-estimation procedure may be used for re-estimating HMM transition probabilities.

Figure 8 illustrates the advantage of employing the constraints of the current invention. This shows a spectral vector $\mathbf{x} = (x_1, x_2, x_3)$, where $|\mathbf{x}| = 1$.

Constraining this spectral vector, e.g. 101 into having a constant modulus has the implication that the class means 102 will all lie on the surface of a hypersphere. In the case shown the hypersphere has two dimensions, and so is an ordinary 2-sphere 103 in an ordinary three-dimensional space.

- 10 Constraining the covariance matrices to be isotropic and diagonal has the effect that the individual classes will project onto this hypersphere in the form of circles 104. This arrangement allows individual speech-sounds to be represented in the spectral-vector space not as individual Gaussian ellipsoids, as is conventional, but as assemblies 105 of many smaller Gaussian
- 15 hypercircles 104 tiling the unit hypersphere 103, offering in the potential for more faithful representation of highly complex-shaped speech-sounds, and thus improved classification performance. Each class (hypercircle) eg 104 will span just a small area within the complex shape that delimits the set of all spectral vectors (which must all lie on the spectral-vector hypersphere 103)
- 20 that could correspond to alternative pronunciations of a particular individual speech-sound; collectively, many such classes 104 will be able to span that whole complex shape much more faithfully than could a single, anisotropic ellipsoid as is conventionally used to represent an individual speech sound. Other sets of Gaussian classes within the same mixture model will be able to
- 25 span parts of other complex shapes on the spectral vector hypersphere, i.e. of other speech sounds. The posterior probabilities associated with each of these Gaussian classes (hypercircles) is a measure of how close the current spectral vector is (on the spectral-vector hypersphere) to the corresponding Gaussian class mean 102 (hypercircle centre). Learning which sets of classes
- 30 correspond to which speech sounds, on the basis of all the temporal correlations between them that are present in the training speech sequences, is the function of the GMM-based HMM, whose inputs are fed from the set of all those posterior probabilities.

To use an analogy, a large number of hypercircles helps one to avoid local minima far better than would a small number of anisotropic ellipsoids, for effectively the same reason that a bunch of sticks gets tangled more easily than a tray of marbles. (In this analogy, minimising the total gravitational potential of the set of marbles plays the analogous role to maximising the model likelihood.) Similarly, one can map out highly complex shapes much more faithfully by using a lot of marbles than by using a few sticks.

The skilled person will be aware that other embodiments within the scope of the invention may be envisaged, and thus the invention should not be limited to the embodiments as herein described.

References

A.R. Webb, *Statistical Pattern Recognition*, Arnold (London), 1999.

B.H. Juang & L.R. Rabiner, *Hidden Markov models for speech recognition*,

Technometrics 33(3), American Statistical Association, 1991.

Claims

1. A signal processing system for processing a plurality of multi-element data encoding vectors, the system:
 - 5 - having means for deriving the data encoding vectors from input signals;
 - being arranged to process the data encoding vectors using a Gaussian Mixture Model (GMM) based Hidden Markov Model (HMM), the GMM based HMM having at least one class mean vector having multiple
 - 10 elements;
 - being arranged to process the elements of the class mean vector(s) by an iterative optimisation procedure;
 - characterised in that the system is also arranged to scale the elements of the class mean vector(s) during the optimisation procedure to provide for the
 - 15 class mean vector(s) to have constant modulus at each iteration, and to normalise the data encoding vectors input to the GMM based HMM.
2. A system as claimed in claim 1 wherein the GMM based HMM has a covariance matrix, the elements of which remain constrained during the
- 20 optimisation procedure such that the matrix is isotropic and diagonal, and the value of the non zero diagonal elements remain constant throughout the optimisation procedure.
3. A system as claimed in claim 1 or claim 2 wherein prior class
- 25 probabilities associated with the GMM based HMM are constrained to be equal, and to remain unchanged throughout the optimisation procedure.
4. A system as claimed in any of the above claims wherein when the elements of the data encoding vectors represent spectral coefficients, the
- 30 normalisation of the data encoding vectors is such that the data encoding vectors have equal moduli.
5. A system as claimed in claim 4 wherein the modulus of each data encoding vector is independent of the overall spectral power in the vector.

6. A system as claimed in claim 4 or claim 5 wherein elements forming spectral coefficients of a data encoding vector are arranged to be individually proportional to the square root of the power in their corresponding spectral
5 band divided by the square root of the overall power contained in spectral bands represented in the vector.

7. A system as claimed in any of claims 4 to 6 wherein the system is arranged to add at least one additional element to each data encoding vector,
10 wherein the added element(s) encode the overall power contained in spectral bands represented in the vector.

8. A system as claimed in claim 7 wherein the system is arranged to add two elements to each data encoding vector to represent the overall power in
15 spectral bands, these two elements arranged such that the sum of their squares is a constant across all data encoding vectors that represent the spectrum of the input signal.

9. A system as claimed in any of claims 1 to 8 wherein the GMM based
20 HMM provides the observation probabilities for a higher level HMM.

10. A system as claimed in any of claims 1 to 9 wherein the derivation of the data encoding vectors from the input signal involves the use of a GMM, whereby this GMM provides the data encoding vectors to the GMM based
25 HMM that comprise elements derived from the GMM's posterior probabilities.

11. A system as claimed in claim 10 wherein elements of the data encoding vectors input from the GMM to the GMM based HMM are proportional to the square root of posterior probabilities of the additional GMM.

30

12. A system as claimed in claim 10 wherein elements of the data encoding vectors input from the GMM to the GMM based HMM are proportional to posterior probabilities of the additional GMM.

13. A system as claimed in any of claims 9 to 12 wherein the constant values for the modulus of each of the class mean vectors may be different at each level.

5 14. A method of processing a signal, the signal comprising a plurality of multi-element data encoding vectors, wherein the data encoding vectors are derived from an analogue or digital input, and where the method employs at least one Gaussian Mixture Model (GMM) based Hidden Markov Model (HMM), the GMM based HMM having at least one class mean vector having
10 multiple elements, and the elements of the class mean vector(s) are optimised in an iterative procedure, characterised in that the elements of the class mean vectors are scaled during the optimisation procedure such that the class mean vectors have a constant modulus at each iteration, and the data encoding vectors input to the GMM based HMM are processed such that they are
15 normalised.

15. A method as claimed in claim 14 wherein a covariance matrix within the GMM based HMM has one or more elements, all of which are constrained during the optimisation procedure such that the matrix is isotropic and
20 diagonal, and the value of its non zero elements remain constant throughout the optimisation procedure.

16. A method as claimed in claim 14 or claim 15 wherein prior class probabilities associated with the GMM based HMM are constrained to be
25 equal, and to remain unchanged throughout the optimisation procedure.

17. A method as claimed in any of claims 14 to 16 wherein when the elements of the data encoding vectors represent spectral coefficients, the data encoding vectors are scaled in a pre-processing stage before being input to
30 the GMM based HMM, such that the moduli of all data encoding vectors are equal.

18. A method as claimed in claim 17 wherein the modulus of each data encoding vector is independent of the overall power in the vector.

19. A method as claimed in claim 17 or claim 18 wherein elements forming spectral coefficients of a data encoding vector are arranged to be individually proportional to the square root of the power in their corresponding spectral band, divided by the square root of the overall power contained in spectral bands represented in the vector.
20. A method as claimed in any of claims 17 to 19 wherein at least one additional element is added to each data encoding vector, wherein the added element(s) encode the overall power contained in spectral bands represented in the vector.
21. A method as claimed in claim 20 wherein two elements are added to each data encoding vector to represent the overall power in spectral bands, these two elements arranged such that the sum of their squares is a constant across all input vectors that represent the spectrum of the input signal.
22. A method as claimed in any of claims 14 to 21 wherein the GMM based HMM provides the observation probabilities for a higher level HMM.
23. A method as claimed in any of claims 14 to 21 wherein the derivation of the data encoding vectors from the input signal involves the use of a GMM, whereby this GMM provides the data encoding vectors to the GMM based HMM that comprise elements derived from the GMM's posterior probabilities.
24. A method as claimed in claim 23 wherein elements of the data encoding vectors input from the GMM to the GMM based HMM are proportional to the square root of posterior probabilities of the additional GMM.
25. A method as claimed in claim 23 wherein elements of the data encoding vectors input from the GMM to the GMM based HMM are proportional to posterior probabilities of the additional GMM.

26. A method as claimed in any of claims 22 to 25 wherein the constant values for the modulus of each of the class mean vectors may be different at each level.
- 5 27. A signal processing system that has been trained according to the method as described in claim any of claims 14 to 26.
28. A computer program designed to run on a computer and arranged to implement a signal processing system for processing one or more multi-
10 element input vectors, the system:
- having means for deriving the data encoding vectors from input signals;
 - being arranged to process the data encoding vectors using a Gaussian Mixture Model (GMM) based Hidden Markov Model (HMM), the
15 GMM based HMM having at least one class mean vector having multiple elements;
 - being arranged to process the elements of the class mean vector(s) by an iterative optimisation procedure;
- 20 characterised in that the system is also arranged to scale the elements of the class mean vector(s) during the optimisation procedure to provide for the class mean vector(s) to have constant modulus at each iteration, and to normalise the data encoding vectors input to the GMM based HMM.
29. A speech recogniser incorporating a signal processing system for
25 processing one or more multi-element input vectors, the recogniser:
- having means for deriving the data encoding vectors from input signals;
 - being arranged to process the data encoding vectors using a Gaussian Mixture Model (GMM) based Hidden Markov Model (HMM), the
30 GMM based HMM having at least one class mean vector having multiple elements;
 - being arranged to process the elements of the class mean vector(s) by an iterative optimisation procedure;


characterised in that the system is also arranged to scale the elements of the class mean vector(s) during the optimisation procedure to provide for the class mean vector(s) to have constant modulus at each iteration, and to normalise the data encoding vectors input to the GMM based HMM.

Abstract

A signal processing system is disclosed which includes a Gaussian Mixture Model (GMM) based Hidden Markov Model (HMM), parameters of which are constrained during the optimisation procedure. Also disclosed is a constraint system applied to input vectors representing the input signal to the system. The invention is particularly, but not exclusively, related to speech recognition systems. The invention reduces the tendency, common in prior art systems, to get caught in local minima associated with highly anisotropic Gaussian components – which reduces the recogniser performance - by employing the constraint system as above whereby the anisotropy of such components may be minimised. The invention also covers a method of processing a signal, and a speech recogniser trained according to the method.

(Figure 3)

Figure 1



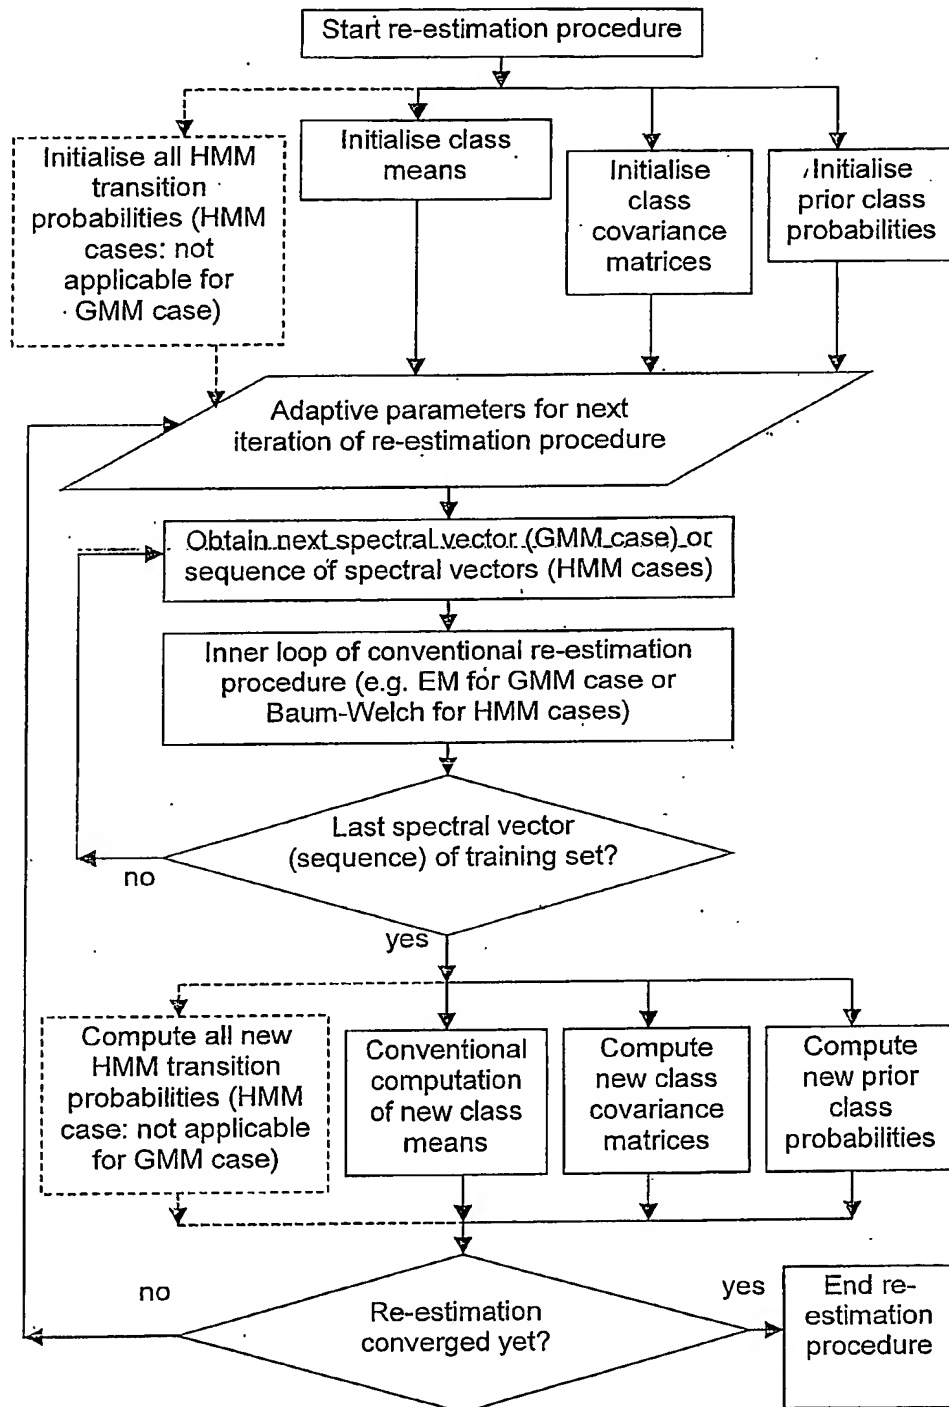


Figure 2

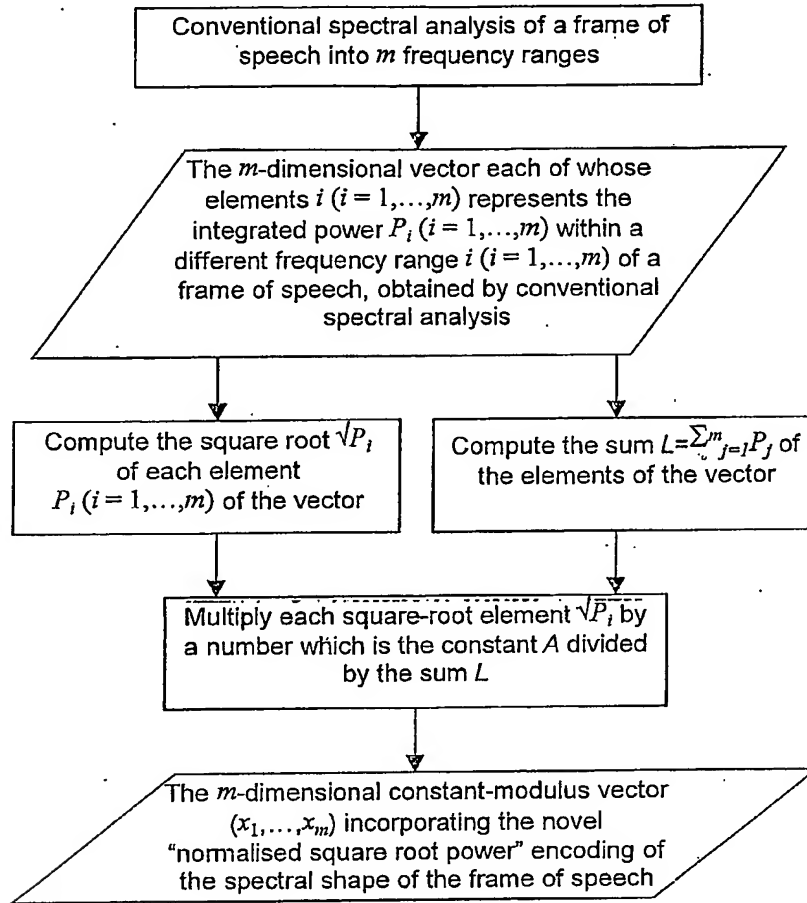


Figure 3

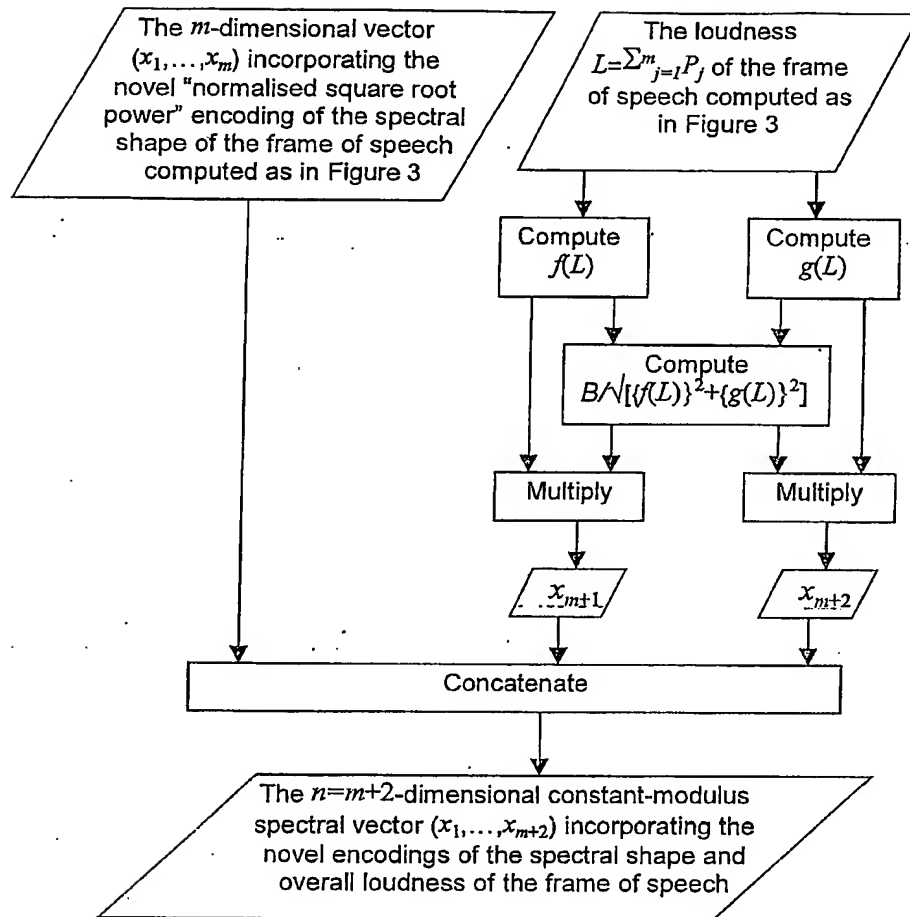


Figure 4

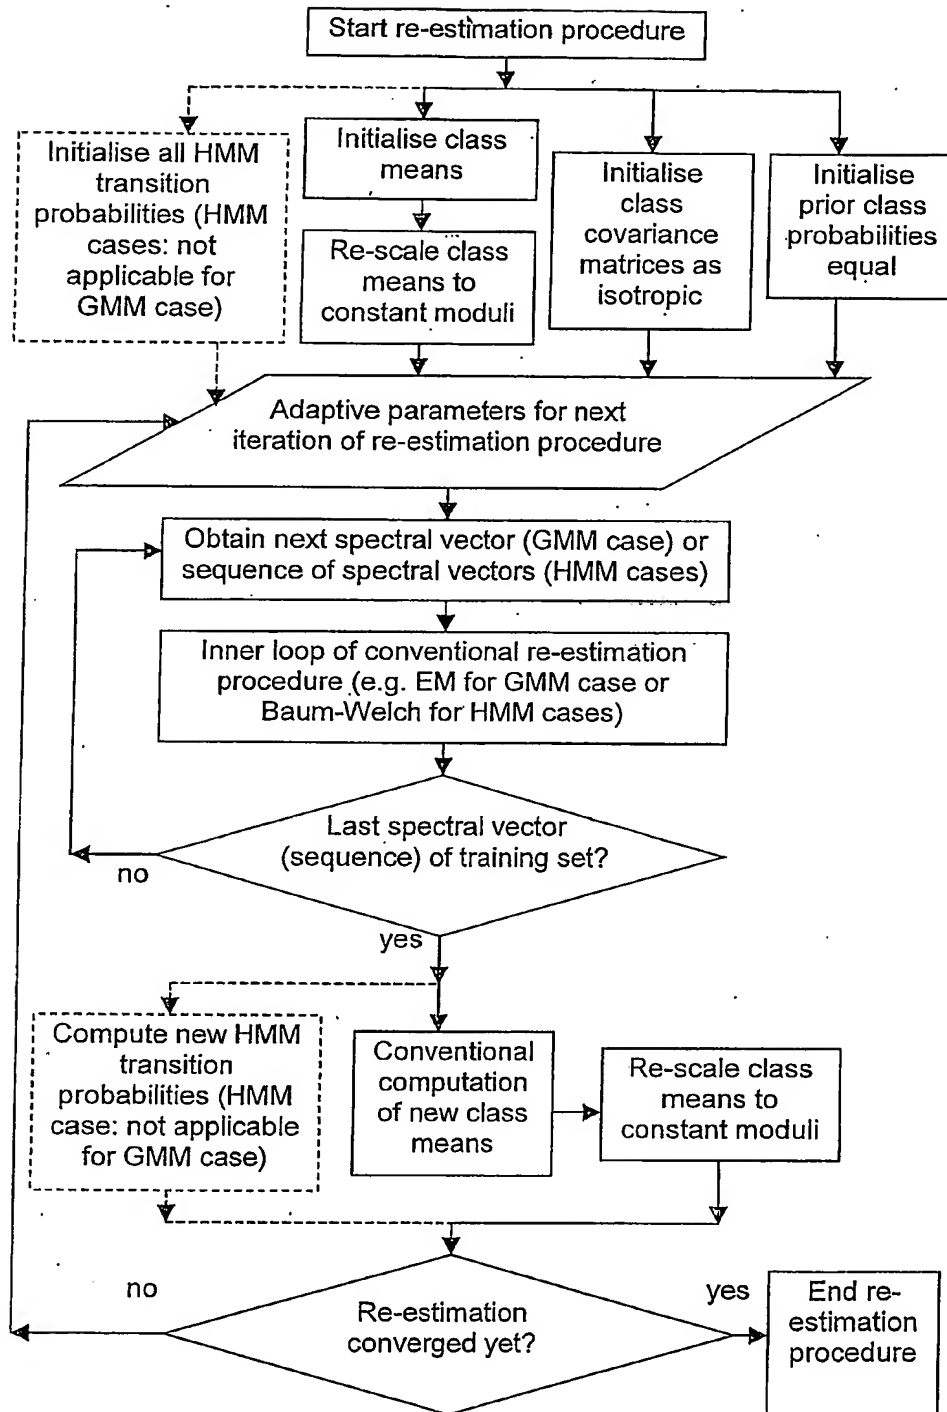


Figure 5

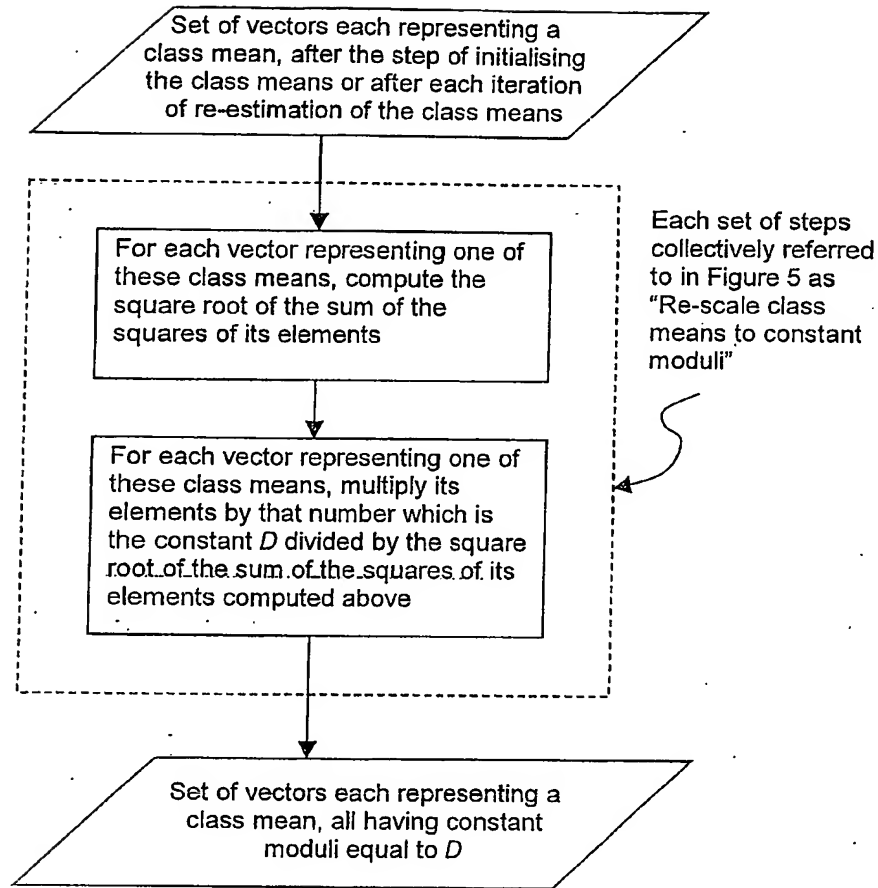


Figure 6

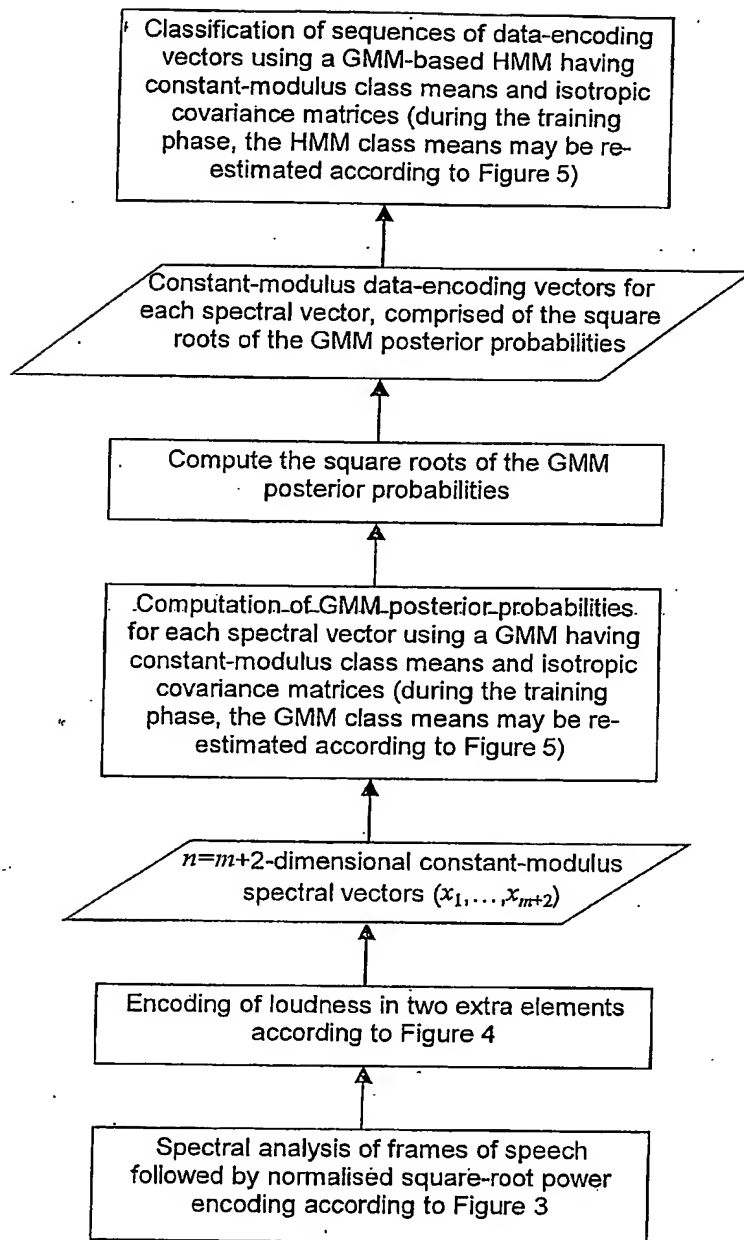



Figure 7

Figure 8



THE PATENT OFFICE

11 APR 2003

Received in Patents
International Unit

Box No. VIII (iv) DECLARATION: INVENTORSHIP (only for the purposes of the designation of the United States of America)
The declaration must conform to the following standardized wording provided for in Section 214; see Notes to Boxes Nos. VIII, VIII (i) to (v) (in general) and the specific Notes to Box No. VIII (iv). If this Box is not used, this sheet should not be included in the request.

**Declaration of inventorship (Rules 4.17(iv) and 51bis.1(a)(iv))
for the purposes of the designation of the United States of America:**

I hereby declare that I believe I am the original, first and sole (if only one inventor is listed below) or joint (if more than one inventor is listed below) inventor of the subject matter which is claimed and for which a patent is sought.

This declaration is directed to the international application of which it forms a part (if filing declaration with application).

This declaration is directed to international application No. PCT/..... (if furnishing declaration pursuant to Rule 26ter).

I hereby declare that my residence, mailing address, and citizenship are as stated next to my name.

I hereby state that I have reviewed and understand the contents of the above-identified international application, including the claims of said application. I have identified in the request of said application, in compliance with PCT Rule 4.10, any claim to foreign priority, and I have identified below, under the heading "Prior Applications," by application number, country or Member of the World Trade Organization, day, month and year of filing, any application for a patent or inventor's certificate filed in a country other than the United States of America, including any PCT international application designating at least one country other than the United States of America, having a filing date before that of the application on which foreign priority is claimed.

Prior Applications: .GB.0207343.5. Filed 28 March.2002.....
Our Ref: IP/P3304.....

I hereby acknowledge the duty to disclose information that is known by me to be material to patentability as defined by 37 C.F.R. § 1.56, including for continuation-in-part applications, material information which became available between the filing date of the prior application and the PCT international filing date of the continuation-in-part application.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Name: Christopher John St Clair WEBBER.....

Residence: United Kingdom.....
(city and either US state, if applicable, or country)

Mailing Address: QinetiQ Malvern, Malvern, St Andrews Road, Bldg EX, Rm 28,
Malvern, Worcs WR14 3PS.....

Citizenship: British.....

Inventor's Signature: C J St Clair Webber..... Date: 21 March, 2003.....
(if not contained in the request, or if declaration is corrected or added under Rule 26ter after the filing of the international application. The signature must be that of the inventor, not that of the agent) (of signature which is not contained in the request, or of the declaration that is corrected or added under Rule 26ter after the filing of the international application)

Name:

Residence:

(city and either US state, if applicable, or country)

Mailing Address:

Citizenship:

Inventor's Signature:

Date:

(if not contained in the request, or if declaration is corrected or added under Rule 26ter after the filing of the international application. The signature must be that of the inventor, not that of the agent) (of signature which is not contained in the request, or of the declaration that is corrected or added under Rule 26ter after the filing of the international application)

☐ This declaration is continued on the following sheet, "Continuation of Box No. VIII (iv)".

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☒ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☒ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.